SEMI-LAGRANGIAN PHYSICS-INFORMED NEURAL NETWORKS (SL-PINNS) FOR SOLVING HYPERBOLIC PARTIAL DIFFERENTIAL EQUATIONS (PDES)

Edoardo Monti, Kateryna Morozovska
Imperial College London
London, UK
{e.monti}@imperial.ac.uk

Aurora Poggi
KTH Royal Institute of Technology
Stockholm, Sweden

Khemraj Shukla
Brown University
Providence, RI, USA

ABSTRACT

Solving nonlinear kinetic systems such as the Vlasov-Poisson partial differential equations (PDEs) is challenging with both classical numerical schemes and vanilla physics-informed neural networks (PINNs) due to the curse of dimensionality, training instabilities in advection-dominated regimes, and prohibitive computational costs for long-time integration. Existing PINN approaches struggle with advection-dominated regimes and require careful constraint tuning to maintain stability over extended time integration. We propose **Semi-Lagrangian PINN (SL-PINN)**, a novel framework that learns the inverse characteristic flow map $\Psi_\theta(t, x, v) \rightarrow (x_0, v_0)$ to reconstruct the distribution as $f(t, x, v) = f_0(\Psi_\theta(t, x, v))$. From a machine learning perspective, the proposed method can be interpreted as learning a transport map that evolves the initial distribution along the characteristic flow of the kinetic equation, while solving the Poisson equation in Eulerian coordinates via spectral methods. Our framework integrates three key innovations: (1) **inverse flow map parameterization**; (2) **hybrid Lagrangian-Eulerian coupling** treating transport and field equations in their natural representations; (3) **volume-preserving loss penalty** that promotes conservation of mass, momentum, and energy.

1 INTRODUCTION

Vlasov-Poisson partial differential equations represent a nonlinear kinetic system that models time evolution of a collisionless plasma consisting of electrons and a self-consistent electrostatic field Birdsall & Langdon (2004); Chen (2016). Essentially they combine the Vlasov equation that models the transport of the electrons and the Poisson equation that models the electrostatic potential Glassey (1996). When solving Vlasov-Poisson PDEs, traditional numerical methods suffer from the complexities associated with higher dimensionality and consequently face higher computational costs Filbet & Sonnendrücker (2003); Crouseilles et al. (2010); Sonnendrücker et al. (1999). At the same time, existing Scientific Machine Learning methods from the family of Physics-Informed Neural Networks (PINNs) Raissi et al. (2019); Karniadakis et al. (2021) experience additional challenges from training instability as well as issues of handling conservation laws Mao et al. (2020); Patel et al. (2022).

Since their appearance in the seminal paper by Raissi et al. (2019), PINNs have gained recognition as a powerful tool for solving forward and inverse problems governed by partial and ordinary differential equations (PDEs and ODEs) and have demonstrated many successful applications to a variety of problems in science and engineering Karniadakis et al. (2021). At the same time, training of PINNs for solving hyperbolic PDEs such as the Vlasov-Poisson system remains fundamentally and persistently challenging - even state-of-the-art variants struggle with this task - since many PINN formulations including vanilla PINNs experience instabilities when solving advection-dominated problems where they suffer from gradient pathologies and convergence failures caused by the hyperbolic nature of transport equations Wang et al. (2021); Krishnapriyan et al. (2021). Additionally, as shown by Rahaman et al. (2019), neural networks (NNs) suffer from spectral bias, hence, for systems with both low- and high-frequency domains, the NN will inherently prioritize learning lower-frequency components while missing the fine-scale structure in the phase space.

These challenges are particularly pronounced for kinetic PDEs, where accurately resolving phase-space structure is critical for capturing collisionless dynamics, plasma instabilities, and long-time behavior in fusion and astrophysical settings. Extending simulations to longer time horizons further exacerbates the difficulty: small approximation errors can accumulate over time, potentially leading to degradation in accuracy or unphysical artifacts Wang et al. (2024a); De Ryck & Mishra (2022).

A central difficulty when training neural solvers for hyperbolic PDEs lies in maintaining physically consistent dynamics. Standard PINN formulations often struggle to control conservation errors, which may manifest as gradual mass drift or artificial dissipation of momentum and energy Jagtap et al. (2022); Mao et al. (2020). These issues are particularly detrimental for kinetic systems, where transport-dominated behavior makes solutions highly sensitive to such structural violations. Some hybrid approaches do attempt to integrate the Semi-Lagrangian methods in the NN learning, but they primarily rely on numerical solvers as an intermediate step Franck et al. (2026).

Several approaches have sought to mitigate these limitations. Causality-respecting PINNs Wang et al. (2024b) modify the training schedule to better propagate information along characteristic directions, improving optimization stability but not directly addressing conservation behavior. Variational formulations Kharazmi et al. (2021) enforce weak solutions that can enhance robustness, though they typically remain dependent on soft penalty terms. Conservative PINNs Jagtap et al. (2020) incorporate discretized conservation constraints, improving invariant preservation at the cost of reduced mesh-free flexibility and increased reliance on predefined grid structures. Energy-based approaches Jin et al. (2021) introduce auxiliary networks to promote Hamiltonian structure, often resulting in substantially higher parameter counts and computational overhead. However, these examples do not fully reconcile structure preservation, robustness in transport-dominated regimes, and practical scalability to higher-dimensional phase spaces, leaving a gap for methods that more naturally align with the geometric properties of kinetic flows.

We propose **SL-PINN** (Semi-Lagrangian PINN), which reformulates the Vlasov-Poisson system by learning the *inverse characteristic flow map* $\Psi_\theta(t, x, v) \rightarrow (x_0, v_0)$ using a neural network. Instead of directly approximating the distribution function $f(x, v, t)$ in Eulerian coordinates, our method parameterizes the inverse flow map that traces particles from their current phase-space location (x, v) at time t to their initial position (x_0, v_0) at $t = 0$. The solution is then reconstructed via $f(t, x, v) = f_0(\Psi_\theta(t, x, v))$, where f_0 is the prescribed initial distribution. This Lagrangian representation exploits the characteristic structure of the Vlasov equation, reducing the hyperbolic transport problem to learning the inverse flow map while solving the Poisson equation on an Eulerian grid via a spectral solver.

Our method combines three key ideas: (1) **Inverse flow-map parameterization:** a neural network learns an approximate backward characteristic map Ψ_θ by enforcing the transport equations along characteristics. This representation preserves positivity and promotes physically consistent phase-space dynamics through a structure-aware formulation. (2) **Hybrid Lagrangian–Eulerian coupling:** the Vlasov transport is modeled in a Lagrangian framework via characteristic tracking, while the Poisson equation is solved on an Eulerian grid, allowing each component to be treated in its natural representation. (3) **Volume-preserving regularization:** a Jacobian-determinant constraint encourages phase-space measure preservation, substantially reducing mass, momentum, and energy drift. By aligning the learning problem with the geometric structure of kinetic flows, SL-PINN improves robustness and long-time accuracy relative to vanilla PINN formulations while retaining the flexibility of neural solvers for complex boundary conditions and inverse problems.

The remainder of the paper is organized as follows. Section 2 introduces the theoretical background of the Vlasov–Poisson system. Section 3 presents the proposed **SL-PINN** framework, including the inverse flow-map formulation, neural network parameterization, hybrid Lagrangian–Eulerian coupling, and the physics-informed training objective with volume-preserving regularization. Section 4 reports numerical experiments on representative benchmarks, such as Gaussian advection and the

Landau damping problem, and analyzes the accuracy and conservation behavior of the proposed method. Finally, Section 5 summarizes the main findings and discusses limitations and directions for future work.

2 BACKGROUND

In this section, the theory underlying this work and the terminology used are briefly reviewed, with emphasis on the Vlassov-Poisson system and characteristic formulation.

2.1 THE VLASOV-POISSON SYSTEM

The Vlasov-Poisson system is a fundamental model in plasma physics and kinetic theory, used to describe the evolution of charged particle distributions in self-consistent electrostatic fields Birdsall & Langdon (2004). The distribution function $f(x, v, t)$ represents particle density at position $x \in \mathbb{R}^d$ with velocity $v \in \mathbb{R}^d$ over time t, evolving according to:

$$\frac{\partial f}{\partial t} + v \cdot \nabla_x f + E(x,t) \cdot \nabla_v f = 0, \tag{1}$$

$$-\Delta\phi = \int_{\mathbb{R}^d} f\, dv - \rho_0, \quad E = -\nabla\phi, \tag{2}$$

where ϕ is the electrostatic potential and ρ_0 is a constant that ensures charge neutrality. This hyperbolic-elliptic coupled system preserves mass, momentum, and total energy:

$$H = \int_x \int_v \frac{1}{2}|v|^2 f dx\, dv + \int_x \frac{1}{2}|\nabla\phi|^2 dx. \tag{3}$$

Computational bottlenecks. The main challenge in solving the Vlasov-Poisson system lies in the curse of dimensionality. The distribution function f lives in a $2d$-dimensional phase space, requiring $O(N^{2d})$ grid points to achieve resolution N per spatial dimension. When considering three-dimensional plasmas ($d = 3$), this yields six-dimensional problems with $\sim 10^{12}$ degrees of freedom at moderate resolution ($N = 100$). Existing methods typically rely on trade-offs: particle-in-cell (PIC) methods reduce complexity to $O(N_p)$ particles but introduce statistical noise Birdsall & Langdon (2004); Eulerian schemes maintain spectral accuracy but scale exponentially Filbet & Sonnendrücker (2003); semi-Lagrangian methods achieve unconditional stability Sonnendrücker et al. (1999) but may violate conservation due to interpolation errors Crouseilles et al. (2010).

2.2 SEMI-LAGRANGIAN METHODS AND CHARACTERISTIC FORMULATIONS

Semi-Lagrangian (SL) schemes solve transport equations by backward-tracking particle trajectories along characteristic curves, yielding unconditional stability (CFL-free time-stepping) Sonnendrücker et al. (1999). The Vlasov equation admits a characteristic formulation: along particle trajectories $(x(t), v(t))$ satisfying the Hamiltonian dynamics:

$$\frac{dx}{dt} = v(t), \tag{4}$$

$$\frac{dv}{dt} = E(x(t), t), \tag{5}$$

$$(x(0), v(0)) = (x_0, v_0), \tag{6}$$

and the distribution function is conserved as:

$$\frac{d}{dt} f(x(t), v(t), t) = 0 \Rightarrow f(x(t), v(t), t) = f_0(x_0, v_0). \tag{7}$$

Let $\Phi_t : (x_0, v_0) \mapsto (x(t), v(t))$ denote the *forward characteristic flow map*. Then the solution can be written as:

$$f(t, x, v) = f_0(\Phi_t^{-1}(x, v)), \tag{8}$$

where Φ_t^{-1} is the *inverse flow map* that is tracing the particles backward, from (x, v) at time t to their initial positions (x_0, v_0).

Classical SL schemes exploit this structure by: (1) tracing particles backward along characteristics to find departure points, (2) interpolating f from a fixed grid at those departure points, and (3) advancing f forward in time without CFL restrictions. However, classical SL methods face two critical limitations for machine learning integration: **(1) Grid-based interpolation errors:** repeated interpolation introduces numerical diffusion and discretization errors, violating conservation laws unless specialized filters are applied Crouseilles et al. (2010); **(2) Lack of adaptivity:** fixed grids cannot adapt to local solution structure, and the method cannot learn from observational data or incorporate data-driven corrections.

3 SEMI-LAGRANGIAN PINNS

To address the limitations of classical SL methods, we reformulate the Vlasov-Poisson system by learning the inverse characteristic flow map. The proposed approach replaces grid-based interpolation with a parametrized neural surrogate that learns the inverse flow map $\Psi_\theta(t, x, v) \approx \Phi_t^{-1}(x, v)$, which traces phase-space coordinates backward along characteristic curves:

$$\Psi_\theta(t, x, v) \mapsto (x_0, v_0).$$

By training Ψ_θ to satisfy the characteristic transport equations via a physics-informed loss, we combine the unconditional stability of semi-Lagrangian schemes with the mesh-free flexibility and data adaptivity of NNs. The solution is thus reconstructed through composition,

$$f_\theta(t, x, v) = f_0(\Psi_\theta(t, x, v)),$$

aligning the neural representation with the geometric structure of kinetic transport. This eliminates interpolation errors while improving the conservation properties inherent to the Lagrangian formulation.

In what follows, we detail the mathematical formulation, neural parameterization, and training procedure underlying the proposed SL-PINN framework.

3.1 CHARACTERISTIC EQUATION

To obtain a governing equation for the inverse flow map, we express the distribution function through the composition $f_\theta(t, x, v) = f_0(\Psi_\theta(t, x, v))$. Substituting this representation into the Vlasov equation yields a partial differential equation for the inverse characteristic map itself. Since the Vlasov equation holds for any admissible initial distribution f_0, the inverse map must satisfy the characteristic transport equation

$$\partial_t \Psi_\theta + v \cdot \nabla_x \Psi_\theta + E(x, t) \cdot \nabla_v \Psi_\theta = 0, \tag{9}$$

$$\Psi_\theta(x, v, t = 0) = (x, v), \tag{10}$$

which expresses the fact that phase-space coordinates remain constant along particle characteristics.

This Lagrangian representation reduces the hyperbolic Vlasov equation to learning the phase-space flow map. By modeling transport through the characteristic mapping rather than directly approximating f, the formulation naturally preserves positivity, mitigates artificial numerical diffusion, and aligns the neural representation with the Hamiltonian structure of the underlying dynamics.

Key properties. This formulation provides several desirable structural properties:

- **Positivity preservation:** If $f_0 \geq 0$, then the reconstructed solution $f(t, x, v) = f_0(\Psi_\theta(t, x, v))$ remains nonnegative.
- **Reduced artificial diffusion:** Modeling transport through characteristic flow maps mitigates the numerical smoothing often observed in Eulerian PINN discretizations.

Additionally, a key observation motivating our formulation is that, when the Vlasov equation is expressed through characteristic dynamics, the transported information is entirely determined by the initial distribution f_0. Consequently, in principle one only needs to resolve the region of the initial phase space (x_0, v_0) where f_0 is non-zero, since outside this support the distribution carries no dynamical information. This observation highlights a potential computational advantage of characteristic-based formulations for high-dimensional cases: the learning process can focus just on the restricted non-trivial portion of phase space that is actually advected by the dynamics.

3.2 SELF-CONSISTENT ELECTRIC FIELD

The electric field $E_\theta(x,t)$ is computed via a hybrid Lagrangian-Eulerian approach. Given the learned f_θ, we compute the charge density by marginalizing over velocity:

$$\rho_\theta(x,t) = \int_{\mathbb{R}} f_\theta(t,x,v)\,dv = \int_{\mathbb{R}} f_0(\Psi_\theta(t,x,v))\,dv, \tag{11}$$

and solve the Poisson equation on an Eulerian grid:

$$\partial_{xx}\phi_\theta(x,t) = \rho_\theta(x,t) - \bar{\rho}, \tag{12}$$

$$E_\theta(x,t) = -\partial_x \phi_\theta(x,t), \tag{13}$$

where $\bar{\rho}$ is the mean charge density ensuring zero-mean potential (for periodic boundaries).

3.3 PHYSICS-INFORMED LOSS

We train f_θ by enforcing consistency with the characteristic transport equation 10 via a physics-informed loss (initial conditions are enforced by construction of the identity map):

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{char}}(\theta) + \lambda_{\text{bc}}\mathcal{L}_{\text{bc}}(\theta) + \lambda_{\text{Jac}}\mathcal{L}_{\text{Jac}}(\theta), \tag{14}$$

where the characteristic consistency loss equation 15 is defined as follows:

$$\mathcal{L}_{\text{char}}(\theta) = \frac{1}{N_c}\sum_{i=1}^{N_c} \left\|\partial_t f_\theta + v_i \nabla_x f_\theta + E_\theta \nabla_v f_\theta\right\|^2, \tag{15}$$

Unlike vanilla PINNs, we do *not* include separate conservation penalty terms (e.g., $\lambda_{\text{mass}}\|\partial_t M\|^2$)—conservation is enforced through the jacobian loss penalty. In fact, to encourage symplectic structure and phase-space volume preservation, as in Hamiltonian dynamics, we penalize deviations of the Jacobian determinant from unity:

$$\mathcal{L}_{\text{Jac}}(\theta) = \frac{1}{N_j}\sum_{k=1}^{N_j} \left(\det \nabla_{(x,v)}\Psi_\theta(t_k, x_k, v_k) - 1\right)^2, \tag{16}$$

Collocation points $\{(t_i, x_i, v_i)\}_{i=1}^{N_c}$ are sampled uniformly from $\Omega = [0,T] \times \Omega_x \times \Omega_v$.

Gradient computation. We compute spatial and velocity gradients $\nabla_x f_\theta$, $\nabla_v f_\theta$ via automatic differentiation. The electric field $E_\theta(x,t)$ is obtained from the Poisson solve at each training iteration. All operations are implemented in PyTorch with gradient computation via autograd.

4 EXPERIMENTS

The proposed SL-PINN model is evaluated on Vlasov-Poisson system and additionally validated using two test cases: advection of a Gaussian distribution and the Landau damping benchmark, a canonical problem with well-established theoretical predictions.

The inverse flow map is approximated with a feedforward neural network as:

$$\Psi_\theta(t,x,v) = \left(x_\theta(t,x,v), x_\theta(t,x,v)\right) \in \mathbb{R}^2, \tag{17}$$

where θ denotes trainable parameters. We use fully-connected NN with input $(t,x,v) \in \mathbb{R}^3$ (for 1D-1V problems). Unless otherwise stated, experiments use a feedforward network with 5 hidden layers of 256 neurons and SiLU activations, trained using AdamW with learning rate $\eta = 10^{-4}$ for 2,000 iterations and 20,480 collocation points per iteration.

4.1 VALIDATION: GAUSSIAN ADVECTION

We first evaluate the proposed method on a smooth advection-dominated Vlasov–Poisson configuration to assess the accuracy of the learned inverse flow map. The computational domain is $x \in [-6, 6]$ (periodic) and $v \in [-6, 6]$, with initial condition

$$f_0(x, v) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(v-u)^2}{2}\right), \qquad u = 0.5,$$

and simulation time $t \in [0, 2]$. A high-resolution semi-Lagrangian solver with FFT-based Poisson solver and Strang splitting is used as a reference solution.

Figure 1 compares the phase-space solution produced by SL-PINN with the reference solver at several time instances. The results show strong agreement between the two solutions, indicating that the learned inverse flow map Ψ_θ accurately reconstructs the transported distribution. Cross-sections of the distribution function along the spatial and velocity directions (Figures 2a and 2b) further confirm that the neural formulation reproduces the correct advection dynamics. The corresponding macroscopic quantities are shown in Figures 3a and 3b, which display the charge density $\rho(x)$ and electric field $E(x)$ at selected times. Both quantities closely match the reference solution throughout the simulation horizon, demonstrating that the flow map representation successfully captures the self-consistent coupling between particle transport and the electrostatic field.

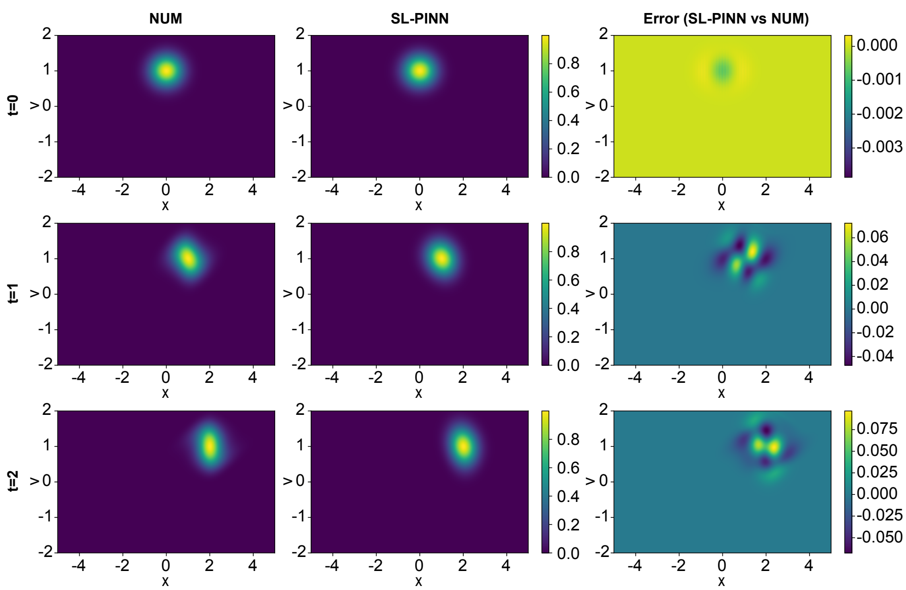

Figure 1: Solution and the error map of the Vlasov-Poisson problem using SL-PINN and the comparison with the high-resolution reference solution at 3 different times.

Finally, characteristic trajectories (Figure 4) illustrate the particle paths reconstructed by the learned inverse map. The close agreement with the reference trajectories confirms that the neural model accurately recovers the underlying characteristic flow.

Comparison with Vanilla PINNs. To highlight the advantages of the proposed formulation, we compare SL-PINN with a vanilla PINN trained using the same neural architecture and optimization settings, but with twice the number of collocation points. Figures 5a and 5b show the evolution of the charge density $\rho(x)$ and electric field $E(x)$ for the two approaches. While SL-PINN closely matches the reference solution, the vanilla PINN formulation struggles to learn the advection dynamics and rapidly flattens the solution after the initial time step.

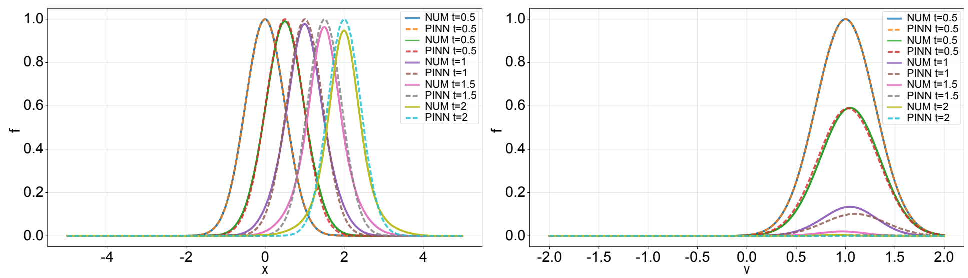

(a) Evolution of f over x (b) Evolution of f over v

Figure 2: Comparative evolution of f over x: (a) and v (b) for numerical solution and SL-PINN solution at $t = 0, 0.5, 1, 1.5, 2$

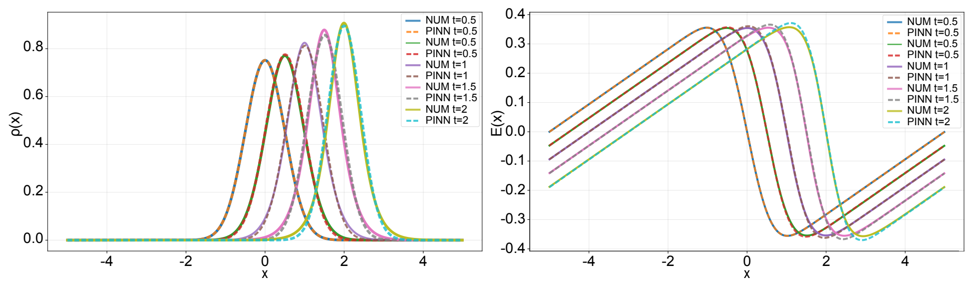

(a) Phase space at $t = 0, 0.5, 1, 1.5, 2$ (b) $E(x)$: SL-PINN vs. theory

Figure 3: Evolution of the corresponding charge density $\rho(x)$: (a) and electric field $E(x)$ (b) at selected time instances

4.2 Validation: Landau Damping

We next consider the Landau damping problem, which tests long-time accuracy and robustness against theoretical predictions. The initial condition is:

$$f_0(x, v) = \frac{1}{\sqrt{2\pi}} e^{-v^2/2} \left(1 + 0.1 \cos x\right), \tag{18}$$

defined on $x \in [-\pi, \pi]$, $v \in [-6, 6]$, and simulated over $t \in [0, 15]$.

A limited hyperparameter sweep was conducted to identify stable configurations for the Landau damping experiment. We varied the learning rate $\eta \in \{10^{-5}, 5 \times 10^{-5}, 10^{-4}\}$, network depth $\in \{5, 7, 10\}$. The sweep reports the validation loss for representative configurations. Several configurations have yielded more stable results as an example for $t \in [0, 10]$ selected setting ($\eta = 10^{-4}$, 5 layers, 256 neurons, $\lambda_{\text{jac}} = 0.001$) yielded stable training and accurate long-time behavior, however, all most stable models have in common demand for longer training above 4000 epochs on AdamW.

In the result representation we use ($\eta = 10^{-4}$, 7 layers, 64 neurons, $\lambda_{\text{jac}} = 0.1$) for $t \in [0, 15]$ over 2000 epochs on AdamW, which already has supplied significantly well performing metrics given the oscillatory nature of the system. Figure 6 shows the phase-space snapshots over x and v as well as the error map. Results show sustained filamentation and fine-scale structure over multiple damping times, indicating robustness in long-time integration, especially if longer training times can be considered. However, fitting of the Electric field and the charge density is not as robust, as shown in Figure 7a and mass preservation is also not complete as indicate by the absolute error for SL-PINN in Figure 7b. We see that with the hyperparameter sweep model can reach significant improvements on the fitting of the Electric field, although some of the spikes persist in time.

4.3 Discussion

The experimental validation demonstrates that SL-PINN successfully learns the inverse characteristic flow map Ψ_θ for the Vlasov–Poisson system. The proposed Lagrangian formulation promotes the

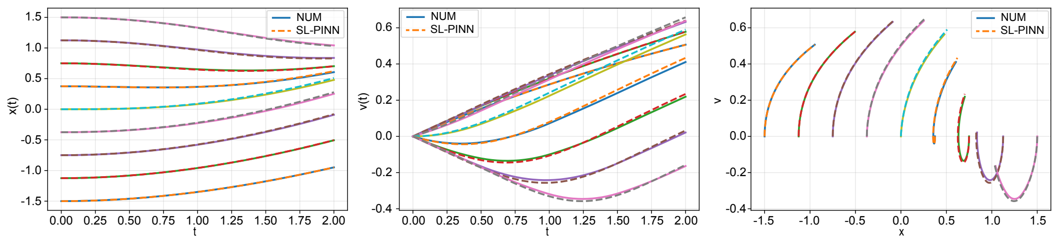

Figure 4: Characteristics $(x(t)$ (on the left) $v(t))$ (central) and phase space $v(x)$ (on the right) for particle tracking.

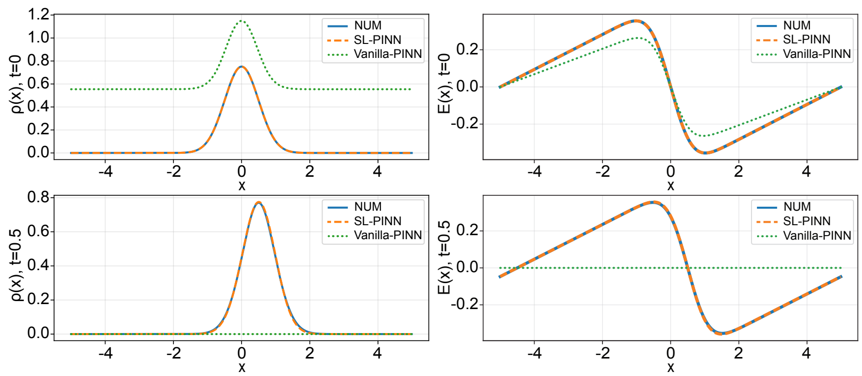

(a) Phase space $\rho(x)$

(b) Electric field $E(x)$

Figure 5: Comparative evolution of the corresponding charge density $\rho(x)$: (a) and electric field $E(x)$ (b) at selected time instances t=0 and t=0.5 for SL-PINN, Vanilla-PINN and numerical solution.

preservation of physical invariants through the structure of the characteristic dynamics together with the Jacobian regularization. Phase-space snapshots (Figures 1, 6) show that the method captures fine-scale filamentation over multiple damping times while exhibiting limited artificial dissipation, indicating that the learned map may accurately reconstruct the underlying transport dynamics.

Despite these encouraging results, several limitations highlight directions for improvement. First, **electric field oscillations** (Figure 7a) are not well-captured by SL-PINN. Improved quadrature schemes, adaptive sampling, or higher-resolution field solvers may help reduce these artifacts.

A promising direction concerns the rigorous enforcement of phase-space volume preservation. In the current formulation this property is encouraged through a soft Jacobian penalty that promotes $\det(\nabla_{x,v}\Psi_\theta) \approx 1$. However, since Vlasov dynamics preserve phase-space volume exactly (Liouville's theorem), it would be desirable to enforce this constraint *by construction*. This could be achieved by parameterizing the inverse map using architectures inspired by *volume-preserving normalizing flows* or recent *flow-matching* formulations, where the neural mapping may be constructed to be symplectic. Such parameterizations would guarantee $\det(\nabla_{x,v}\Psi_\theta) = 1$ exactly, eliminating the need for the Jacobian penalty and potentially improving both training stability and long-time conservation behavior.

5 CONCLUSION

We introduced SL-PINNs, a neural framework for solving the Vlasov–Poisson system by learning the inverse characteristic flow map $\Psi_\theta(t, x, v) \mapsto (x_0, v_0)$. Instead of approximating the distribution function directly, the proposed method reconstructs the solution through composition with the initial distribution, aligning the neural representation with the characteristic structure of kinetic transport. The approach combines three key components: (i) inverse flow-map parameterization, which transforms the hyperbolic transport problem into learning a smooth backward map; (ii) hy-

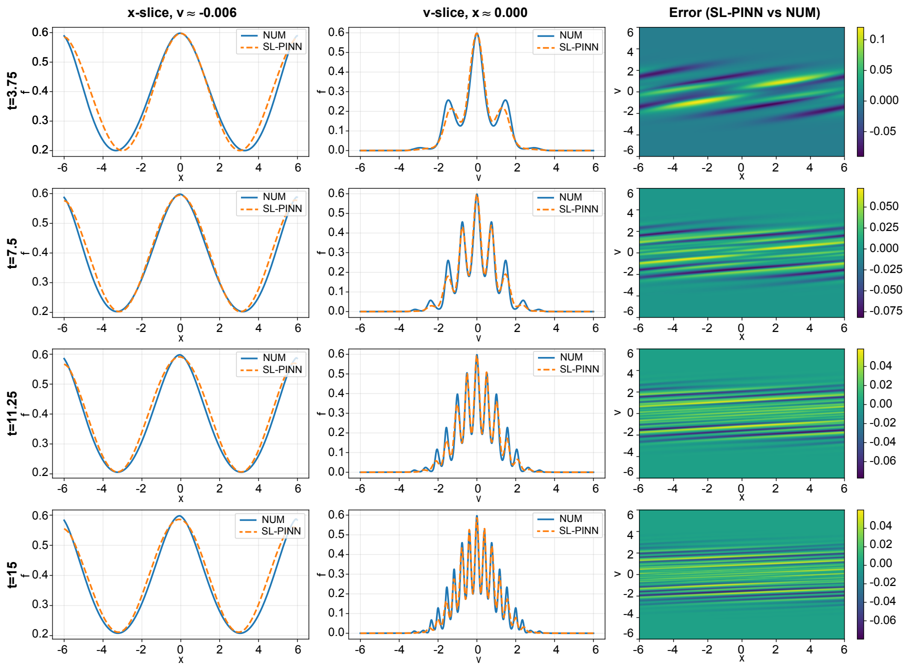

Figure 6: Solution and the error map of the Landau Damping test problem using SL-PINN and the comparison with the high-resolution reference solution at 4 different times $t = 3.75, 7.5, 11.25, 15$.

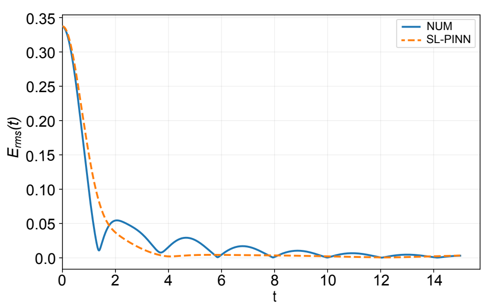

(a) $E_{\text{rms}}(t)$ comparison over time

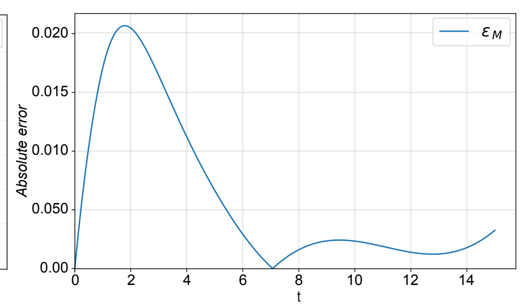

(b) Absolute error for preservation of mass

Figure 7: Energy and mass preservation for Landau Damping problem with $E_{\text{rms}}(t)$: comparison between SL-PINN and numerical solver on the left and the absolute error for mass preservation on the right.

brid Lagrangian–Eulerian coupling, which models particle transport in characteristic coordinates while solving the Poisson equation in Eulerian form; and (iii) volume-preserving regularization, which promotes phase-space measure conservation through a Jacobian constraint.

Experiments on canonical benchmarks, including Gaussian advection and Landau damping, demonstrate that the proposed formulation captures characteristic transport dynamics while preserving fine-scale phase-space structure with limited artificial diffusion. At the same time, the experiments highlight several challenges, including electric-field oscillations caused by finite grid resolution and sensitivity to hyperparameter selection.

In conclusion, the proposed formulation suggests that learning characteristic transport maps provides a promising direction for solving advection-dominated PDEs with neural networks. Future work will explore structure-preserving neural architectures that enforce phase-space volume preservation by construction, as well as extensions to higher-dimensional kinetic systems and other PDEs.

REFERENCES

Charles K Birdsall and A Bruce Langdon. *Plasma physics via computer simulation*. CRC press, 2004.

Francis F Chen. *Introduction to plasma physics and controlled fusion*, volume 1. Springer, 2016.

Nicolas Crouseilles, Michel Mehrenberger, and Eric Sonnendrücker. Conservative semi-lagrangian schemes for vlasov equations. *Journal of Computational Physics*, 229(6):1927–1953, 2010.

Tim De Ryck and Siddhartha Mishra. Generic bounds on the approximation error for physics-informed (and) operator learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 10945–10958, 2022.

Francis Filbet and Eric Sonnendrücker. Comparison of eulerian vlasov solvers. *Computer Physics Communications*, 150(3):247–266, 2003.

Emmanuel Franck, Victor Michel-Dansac, Laurent Navoret, and Vincent Vigon. Neural semi-lagrangian method for high-dimensional advection-diffusion problems. *Computer Methods in Applied Mechanics and Engineering*, 448:118481, 2026.

Robert T Glassey. *The Cauchy problem in kinetic theory*. Society for Industrial and Applied Mathematics, 1996.

Ameya D Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 365:113028, 2020. doi: 10.1016/j.cma.2020.113028.

Ameya D Jagtap, Zhiping Mao, Nikolaus Adams, and George Em Karniadakis. Physics-informed neural networks for inverse problems in supersonic flows. *Journal of Computational Physics*, 466:111402, 2022.

Xiaowei Jin, Shengze Cai, Hui Li, and George Em Karniadakis. Nsfnets (navier-stokes flow nets): Physics-informed neural networks for the incompressible navier-stokes equations. *Journal of Computational Physics*, 426:109951, 2021. doi: 10.1016/j.jcp.2020.109951.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Ehsan Kharazmi, Zhongqiang Zhang, and George Em Karniadakis. hp-vpinns: Variational physics-informed neural networks with domain decomposition. *Computer Methods in Applied Mechanics and Engineering*, 374:113547, 2021. doi: 10.1016/j.cma.2020.113547.

Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pp. 26548–26560, 2021.

Zhiping Mao, Ameya D Jagtap, and George Em Karniadakis. Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering*, 360:112789, 2020.

Ravi G Patel, Indu Manickam, Nathaniel A Trask, Mitchell A Wood, Myunghyun Lee, Ignacio Tomas, and Eric C Cyr. Thermodynamically consistent physics-informed neural networks for hyperbolic systems. *Journal of Computational Physics*, 449:110754, 2022.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

Eric Sonnendrücker, Jean Roche, Pierre Bertrand, and Alain Ghizzo. The semi-lagrangian method for the numerical resolution of the vlasov equation. *Journal of Computational Physics*, 149(2): 201–220, 1999.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421: 116813, 2024a. doi: 10.1016/j.cma.2024.116813. Preprint originally released in 2021 as arXiv:2203.07404.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421: 116813, 2024b.